(12) UK Patent Application (19) GB (11) 2 086 625 A

(21) Application No 8132560
(22) Date of filing
29 Oct 1981
(30) Priority data
(31) 203412
(32) 3 Nov 1980
(33) United States of America (US)
(43) Application published
12 May 1982
(51) INT CL³ G06F 3/00
(52) Domestic classification
G4A 17B FG
(56) Documents cited
GB 2013008A
(58) Field of search
G4A

(71) Applicant
Western Electric Company Incorporated
222 Broadway
New York
NY 10038
United States of America

(72) Inventor
Radhakrishna Shastri Divakaruni
(74) Agents
B R Lawrence
Western Electric Company Limited
5 Mornington Road
Woodford Green
Essex IG8 0TU

(54) Disc intercommunication system

(57) A disc intercommunication system is disclosed for use in a data processing system and provides a communication path (218) which interconnects all the programmable disc devices (212–214) in a multiprocessor system. The disc devices (212–214) use their self-contained processor (220) and memory (221) to interpret data transfer commands and thereby determine which disc devices are to be involved in the data file transfer. The disc devices (212–214) communicate among themselves directly via the communication path (218) to accomplish the data file transfer thereby saving considerable overhead in the host processors.

FIG. 2

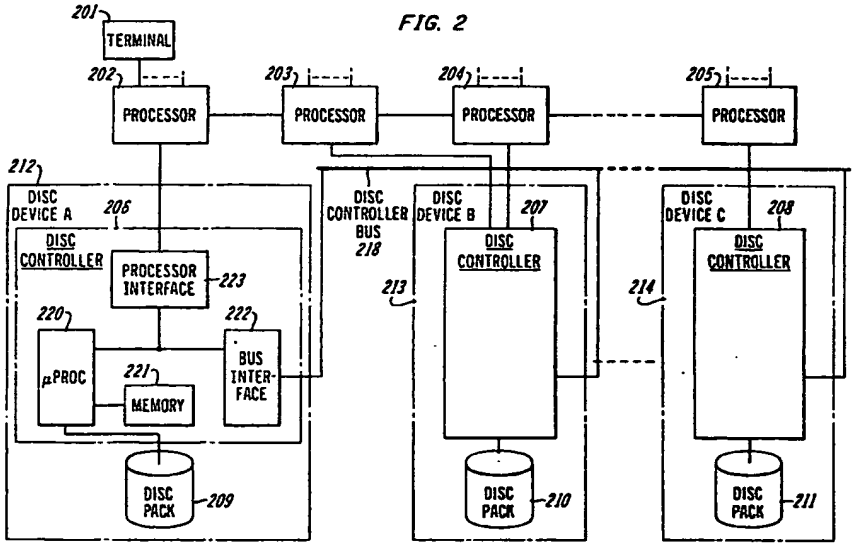

GB 2086625 A

FIG. 1
PRIOR ART

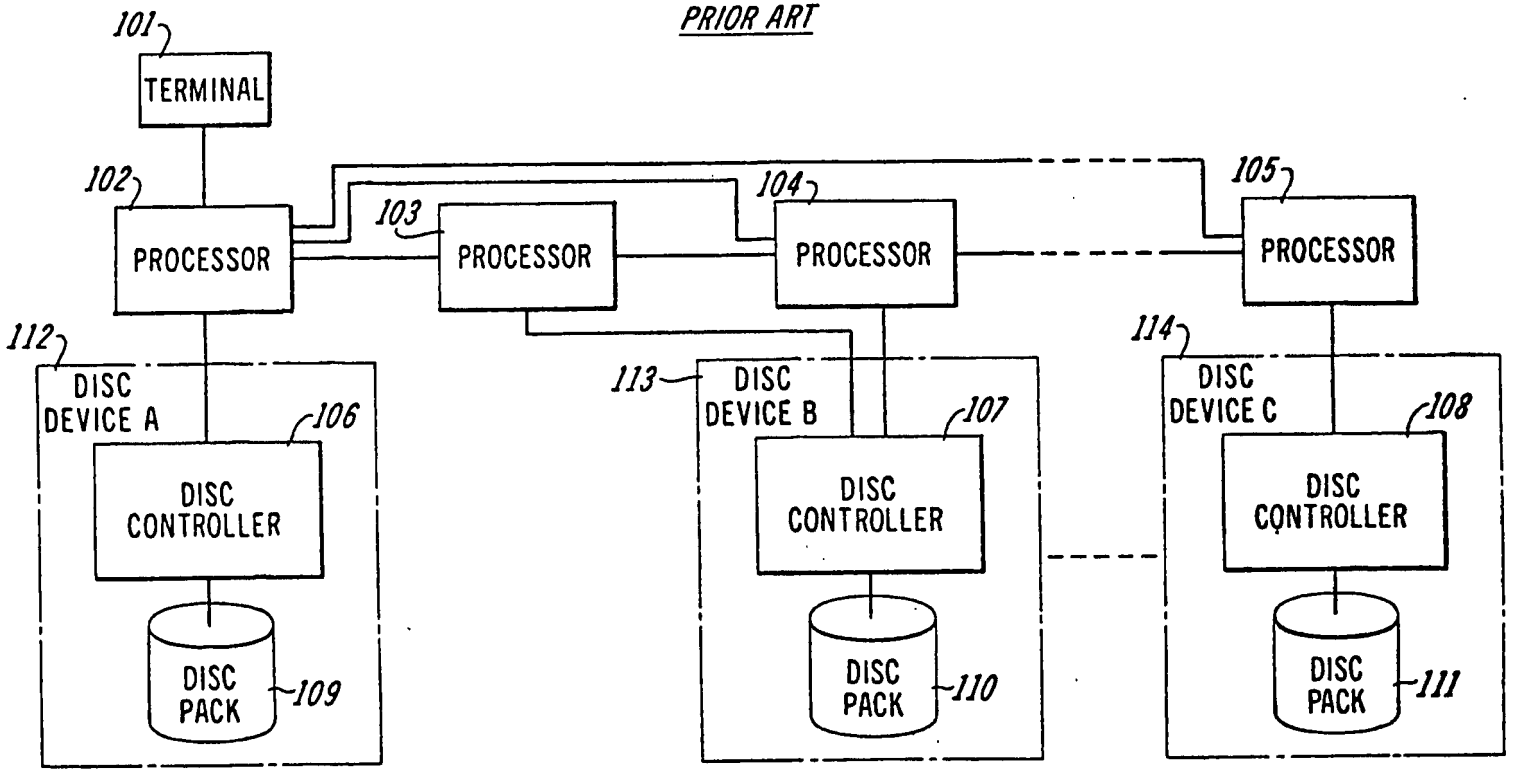

FIG. 3

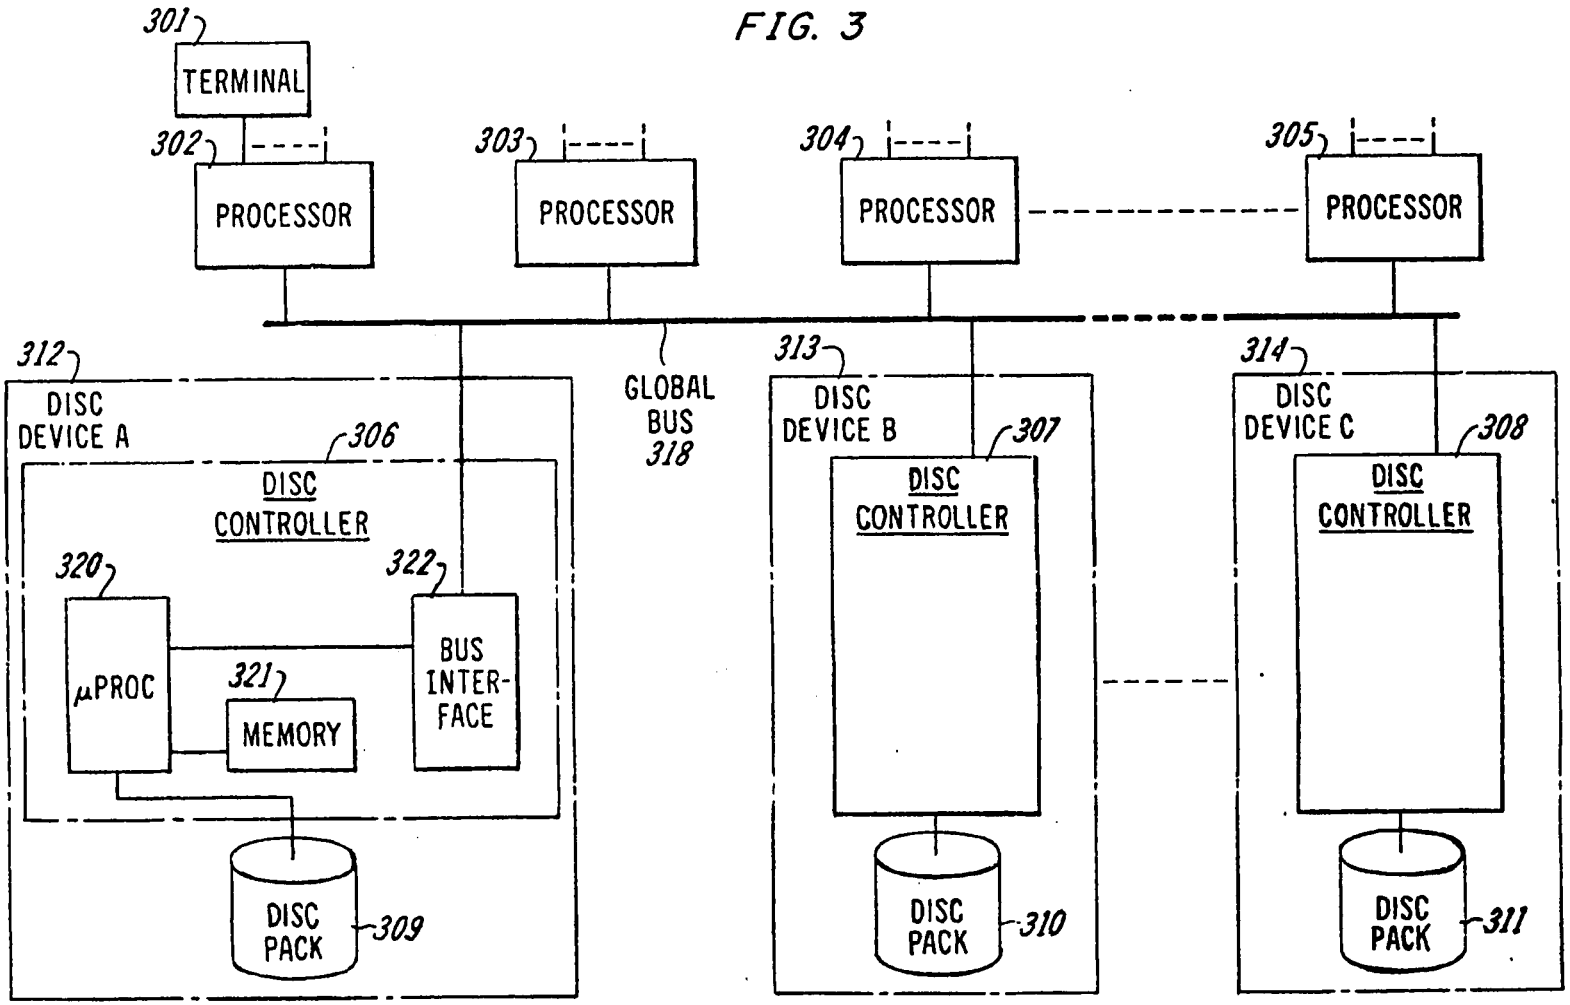
301
TERMINAL
302
PROCESSOR
303
PROCESSOR
304
PROCESSOR
305
PROCESSOR
GLOBAL BUS 318
312
DISC DEVICE A
306
DISC CONTROLLER
320
μPROC
321
MEMORY
322
BUS INTER-FACE
DISC PACK
309
313
DISC DEVICE B
307
DISC CONTROLLER
DISC PACK
310
314
DISC DEVICE C
308
DISC CONTROLLER
DISC PACK
311

SPECIFICATION

Disc intercommunication system

This invention relates to a disc intercommunication sysem for use in a data processing system.

Traditionally, the disc devices used in data processing systems contain the storage media which is the disc pack with its associated read/write heads and drive mechanism, and a simple controller. The disc packs are "dumb" devices, having no built-in intelligence. As a result, a separate controller is required for each disc pack to supply the operating logic. These controllers are expensive and only one disc pack can be operated at a time, thereby rendering all other disc packs connected to the same controller effectively inaccessible for the processor while the one disc pack is being accessed. In particular, the processor, in accessing a disc device, makes a request to the selected disc device to read or write. This request contains a precise address identifying a particular location on the storage media (disc pack). The controller in the selected disc device receives this address, activates the appropriate read/write head and performs the requested operation. During this period, the requesting processor is waiting for the result from the controller, an operation which is termed: wait-I/O operation. This period of inactivity is a waste of processor real time and the processor can therefore be heavily I/O bound, which limits the capacity of the processor.

A variation of this traditional disc device structure is the disc device which contains a built-in intelligent controller employing a microprocessor, such as the STC 2700 series of disc devices. In these units, the operating intelligence is internal to the disc device and the disc controller can perform some of the functions normally assumed by the associated processor. The STC 2700 series disc device is described in the article by Mr. Pranger entitled "'Intelligent Disc Drive for the 1980's", which appeared in the February, 1979 issue of Mini Micro Systems on pages 72–78.

According to the present invention there is provided a disc intercommunication system for use in a data processing system which includes one or more processors, said disc intercommunication system comprising a plurality of disc drives, a plurality of control means for interfacing the disc drives with said one or more processors, and means interconnecting the control means for transferring data relating to said disc drives between said control means.

In a preferred system in accordance with the present invention, the means interconnecting the control means will comprise bus means connected to all of said control means.

Some exemplary embodiments of the invention will now be described reference being made to the accompanying drawings, in which:

Figure 1 shows a typical prior art multiprocessor-disc interconnection scheme;

Figure 2 depicts the details of a disc interconnection scheme system in accordance with the present invention; and

Figure 3 illustrates a variation of the system of Fig. 2.

A typical prior art multiprocessor computer system is illustrated in Fig. 1. This system comprises a plurality of processors 102–105 which are interconnected in well-known fashion to serve a plurality of terminals, I/O ports, etc. Each of processors 102–105 is connected to a number of these devices although, for simplicity of illustration, only terminal 101 is shown in Fig. 1.

In addition to the aforementioned devices, the typical prior art multiprocessor system includes a plurality of disc devices 112–114 which are connected to processors 102–105 in any of a number of well-known configurations. The system of Fig. 1 illustrates two such typical connections: a processor 102 connected to a dedicated disc device 112 and a disc device 113 which is shared by two processors (103 and 104). Each of the disc devices 112–114 is itself comprised of the storage media (disc pack 109–111 respectively) with its associated drive mechanism and read/write heads, as well as a controller (106–108 respectively) which supplies the operating logic for the associated disc pack and which interfaces the associated processor with the disc pack. As can be seen from this configuration, all interdisc communication requires the involvement of two processors as does the accessing of a disc device by a processor not directly connected to said disc device, since there is no direct interconnection of the disc devices.

Two operations will be used throughout this text as illustrations: simple data access and data update of multiple file copies. Using the system of Fig. 1, assume in the first case that the operator of terminal 101 requests a file which resides in data disc pack 110 of disc device 113. Terminal 101 transmits the request to processor 102 which searches its memory to determine the location and size of the requested file. This information is formatted by processor 102 into a data message, which message contains several elements, which are typically: file identifier (name), file address, file size. Having discovered the file's location (in disc pack 110 of disc device 113) and having generated the data message, processor 102 relays the data message to processor 104 (via their multiprocessor interconnection as shown in Fig. 1), which processor translates this data message into a hardware disc address for access purposes, This hardware disc address is then transmitted

by processor 104 to disc controller 107 of disc device 113. Disc controller 107 uses this hardware address information to operate the read heads and other necessary mechanisms of disc pack 110 to read the requested file from disc pack 110 and to transmit this file to processor 104. Then the data file, accessed through associated disc controller 107 of disc device 113 by processor 104, is transmitted by processor 104 directly to processor 102 (via their multiprocessor interconnection as shown in Fig. 1), which processor 102 supplies the file (in well-known fashion) to the user at terminal 101.

As another example of the operation of the system of Fig. 1, consider how data is updated in multiple copies of a file. Having accessed and edited a file, the operator of terminal 101 would send the updated file to processor 102 for storage. A search of its memory reveals to processor 102 that the file to be replaced is stored in two locations (say disc pack 110 of disc device 113 and disc pack 109 of disc device 112) for reliability or operational reasons. Processor 102 uses its associated disc controller 106 of disc device 112 to update the file in disc pack 109 in well-known fashion, and uses processor 104 and its associated disc controller 107 of disc device 113 to update the file in disc pack 110. However, in performing these operations, processor 102 must once again determine the size of the file and the locations where it is to be stored. This information is formatted by processor 102 into a data message, which message contains several elements, which are typically: file identifier (name), file address, file size. Processor 102 transmits this data message along with the file to be stored directly to processor 104 via the interconnection shown between processors in this multiprocessor system. Processor 104 translates this data message into a hardware disc address for access purposes and transmits this hardware disc address and the file to disc controller 107 of disc device 113. Disc controller 107 uses this hardware address information to operate the write heads and other necessary mechanisms of disc pack 110 to thereby write the file into disc pack 110. Similarly, processor 102 would generate a hardware disc address for disc controller 106 of disc device 112 and transmit this hardware disc address and the file to disc controller 106 so that disc controller 106 will write the file into disc pack 109.

Thus, in the system of Fig. 1, if the disc device to be accessed is not directly associated with the controlling processor, the processor associated with the disc device to be accessed is required to participate in the data transfer. This need for the involvement of additional processors in a data transfer is a waste of processor real time, since the processor functions merely as a buffer and address generator in these situations, interfacing the selected disc device with the requesting processor.

In the subject invention, processors, terminals and disc devices are interconnected as shown in Fig. 2. Like Fig. 1, this drawing illustrates several processors which are interconnected in well-known fashion to serve a plurality of terminals, I/O ports etc. Each of processors 202–205 is connected to a number of these devices although, as in Fig. 1, only one terminal 201 is illustrated. In addition, a plurality of disc devices 212–214 are connected to processors 202–205 as in Fig. 1 and each of these disc devices 212–214 is comprised of the storage media (disc pack 209–211 respectively) with its associated drive mechanism and read write heads as well as a controller (206–208 respectively). Unlike Fig. 1, however, Fig. 2 illustrates a network of disc controllers 206–208 interconnected by disc controller bus 218. Additionally, each disc controller 206–208 is of the type that contains built-in intelligence. In particular, a typical controller 206 contains a microprocessor 220, memory 221, as well as processor interface 223 to couple microprocessor 220 to processor 202, and bus interface 222 to couple microprocessor 220 to disc controller bus 218.

The same examples will be used in the description of Fig. 2 as were used to describe the prior art shown in Fig. 1, and differences in operation between the two methods will be amplified after this description. Assume again that the operator of terminal 201 requests data stored in disc pack 210 of disc device 213. Terminal 201 transmits the request to processor 202 and this request is immediately transferred, usually without modification, by processor 202 to disc controller 206 of disc device 212. Disc controller 206 is an intelligent controller which may be a separate unit or a controller internal to a disc device such as that described in the article by M. Pranger entitled "Intelligent Disc Drive for the 1980's", which appeared in the February, 1979 issue of Mini Micro Systems on pages 72–78. The disc controller 206 shown in Fig. 2 itself contains a relatively small processor 220 (a microprocessor) and processor interface 223 receives the request from processor 202 and applies the request to microprocessor 220. Microprocessor 220 receives this request and operates under control of a program stored in microprocessor 220 and/or memory 221 to search the disc controller memory 221 (or possibly the disc pack 209 itself) to determine the location and size of the requested file. This information is formatted by microprocessor 220 into a data message, which message contains several elements, which are typically: file identifier (name), file address, file size. Having found the file's address in disc pack 210, disc controller 206

relays the request for data directly to disc controller 207 of disc device 213 via disc controller bus 218. This is accomplished by microprocessor 220 transmitting the data message that it generated, via bus interface 222, to disc controller bus 218. Since all the disc controllers 206–208 are connected to disc controller bus 218, this data message will reach its signified destination, disc controller 207, in well-known fashion.

Disc controller 207 also contain the elements (220–223) which are shown in Fig. 2 for disc controller 206. Therefore, the data message placed on disc controller bus 218 by microprocessor 220 will be received by the bus interface and thence the microprocessor in disc controller 207. This microprocessor will, as described above for the prior art system, translate this data message into a hardware disc address and use this information to operate the read heads and other necessary mechanisms of disc pack 210 to read the requested file from disc pack 210. The microprocessor in disc controller 207 then transmits the aforementioned data message with the requested file to disc controller bus 218 via the bus interface in disc controller 207. Microprocessor 220 of disc controller 206 monitors disc controller bus 218 in well-known fashion, and thereby detects the presence of this data message placed on bus 218 by disc controller 207. Microprocessor 220 receives the data message and the requested file from disc pack 210 via bus interface 220. Then the requested file, accessed in disc pack 210 through disc controller 207 by disc controller 206, is transmitted by microprocessor 220 via processor interface 223 to processor 202 for execution.

In the second illustrative example, data is to be updated in redundant files. Having accessed and edited a file, the operator of terminal 201 might send the updated file for storage to processor 202, which processor immediately relays the file and a storage command to disc controller 206 of disc device 212. Assume that a memory search revealed to disc controller 206 that the file to be replaced was stored in two separate disc packs (say 209 and 210) for reliability or operational reasons. If so, disc controller 206 directly updates the file in disc pack 209 in standard fashion, and uses disc controller 207 (via disc controller bus 218) to update the file in disc pack 210.

This is accomplished by microprocessor 220 receiving the request and the file from processor 202 via processor interface 223. Microprocessor 220 searches the disc controller memory 221 to determine the destinations of the file. This information is formatted by microprocessor 220 into a data message, which message contains several elements which are typically: file identifier (name), file address, file size. Having found the file's address in disc pack 210, microprocessor 220 outputs the data message and the file to disc controller bus 218 via bus interface 222. Since all the controllers 206–208 are connected to disc controller bus 218, this data message and file will reach their signified, destination, disc controller 207, in well-known fashion.

Disc controller 207, as mentioned above, also contains the elements (220–223) which are shown in Fig. 2 for disc controller 206. Therefore, the data message will be received by the microprocessor, via the bus interface, of disc controller 207. This microprocessor will translate this data message into a hardware address and use this information to operate the write heads and other necessary mechanisms of disc pack 210 to write the received file into disc pack 210.

Similarly, microprocessor 220 would generate a hardware address and use this information to operate the write heads and other necessary mechanisms of disc pack 209 to also write the file into disc pack 209.

Differences

Comparing Fig. 1 to Fig. 2 reveals that the significant physical difference in overall structure between the prior art and the subject system is that the disc controllers of the subject system are interconnected by a disc controller bus. This leads to several differences in operation. In the first example given, processor 102 (Fig. 1) of the prior art system receives a data request and then searches its memory for the file location before relaying the data request, now in the form of a data message, to processor 104, which processor is associated with the requested disc device 113. In contrast, processor 202 (Fig. 2) of the subject invention simply relays the received request without modification to disc controller 206 and depends on disc controller 206 to determine the file address. In the system of Fig. 1, the file address, having been relayed to processor 104, causes processor 104 to generate a hardware address for the requested file and then read the file from disc pack 110 via disc controller 107. Processor 104, by way of processor 102, then transfers this file to the terminal (101). In the subject system, however, no processor was occupied with the file transfer since disc controller 206 found the file address, accessed the file via disc controller bus 218 and disc controller 207, and presented the complete file to processor 202 for transfer to the terminal (201).

In the prior art, therefore, processor real time was used to access data, specifically: time was taken by processor 102 to find a file location in its mapping memory, then generate and transmit a data message to processor 104. Also, time was taken by processor 104 to translate the data message to an absolute disc address, time was taken by processor

104 to access the data, and time was taken by processor 104 to transfer the data to processor 102. In the subject invention, however, these tasks are all undertaken by the several controllers (206 & 207 in the example given) so that processors 202 through 205 may be free for other work.

In the second illustrative example given, processor 101 in the prior art structure again started its task by searching its memory for the locations of the file to be updated before relaying the write command. Then processor 102 wrote data into duplicate files by way of disc controller 106, processor 104, and disc controller 107. Processor 202 of the subject system simply relays the write command to disc controller 206, which device (206) controls the entire writing task for the requesting processor (202).

Again in this example, processor real time was used in the prior art system to write data, specifically: time was taken by processor 102 to find locations in its mapping memory, then generate and transmit a data message to processor 104. Also, time was taken by processor 104 to convert this data message into absolute disc address, time was taken by both processors 102 and 104 to transfer data between them, and time was taken by processor 104 to transfer data to disc controller 107. And again, in the subject invention, all these tasks are undertaken by the several controllers (206 & 207 in the example given) so that processors 202 through 205 may be free for other work.

An additional advantage of the subject invention is that, with this structure, it is possible to have a single file spanning several disc packs and this fact would be logically transparent to the system user. This, therefore, provides the capability for having file sizes larger than the capacity of a single disc pack. The intelligent controller in the disc device where the file begins would have the necessary file header information and would manage the division and distribution of the file segments to other disc devices. In fact, in general, any file located on any disc device can be logically accessed from any processor.

The concept of disc controller intercommunication may be realized in many ways. For example, Fig. 2 illustrates interconnection of the disc controllers by means of a data bus: disc controller bus 218. Another example of disc interconnection is the global bus structure shown in Fig. 3 wherein there is total interconnection between processors and controllers. The advantage of such a system (as compared with the other embodiment) is that communication between processors and controllers is more direct and hence faster. Its disadvantage is that only one message at a time can be transmitted between component assemblies. For example, processor 304 cannot send a message to processor 305 while processor 302 is sending data to disc controller 306 via global bus 318. The operation of this system is essentially the same as that of the system of Fig. 2 with the exception that each processor would be capable of communicating with all disc devices directly and could therefore directly receive files via global bus 318, transmitted by a disc device not associated with that processor.

Obviously, since there is a global bus 318 interconnecting all processors 302–305 and disc devices 312–314, the internal system structure must differ from that of the system of Fig. 2. This difference is that disc controllers 306–308 do not require a processor interface 223 since all extra-disc communication is via global bus 318 and thus bus interface 322 is the universal communication interface for disc controller 306. In operation, typically, each processor (e.g. 302) is assigned a disc device (e.g. 312) and communication between them is by data messages transmitted via global bus 318. Microprocessor 320 monitors global bus 318 via bus interface 322 and recognizes a data message from processor 302 in well-known fashion rather than receiving messages from processor 302 via a dedicated communication path (as in Fig. 2). Otherwise, the operational description recited above for Fig. 2 is directly applicable to the system of Fig. 3, with the added advantage that the disc device 313 having the requested file can transmit the file either directly to the requesting processor 302 or to the disc device 312 which is associated with that processor.

While a specific embodiment of the invention has been disclosed, variations in structural detail, within the scope of the appended claims, are possible and are contemplated.

CLAIMS

1. A disc intercommunication system for use in a data processing system which includes one or more processors, said disc intercommunication system comprising a plurality of disc drives, a plurality of control means for interfacing the disc drives with said one or more processors, and means interconnecting the control means for transferring data relating to said disc drives between said control means.

2. A system as claimed in claim 1, in which the means interconnecting the control means comprises bus means connected to all of said control means.

3. A system as claimed in claim 2, in which each of the control means is responsive to a transfer request signal by one of said one or more processors for interpreting the data transfer request and for directing one or more of the plurality of disc drives to transfer the requested data to the requesting processor via the bus means.

4. A system as claimed in claim 3, in which the bus means is connected to and

interconnectes all of said one or more processors with all of said control means.

5. A system as claimed in claim 3, in which each of the control means comprises a disc controller which includes a processor means which in response to a data transfer request by one of said one or more processors identifies the disc drive containing the data, and a bus interface means which in response to commands from the processor means transmits a request onto the bus means, the bus interface means associated with the disc drive containing the data being effective for transferring the data to the processor requesting the data.

6. A system as claimed in claim 5, in which the disc controller comprises a memory, the processor means being responsive to the data transfer request for acessing the memory to obtain location data for its data and for generating a data message identifying the location, and the bus interface means being responsive to the data message for outputting the data message to the data bus means.

7. A system as claimed in claim 5, in which the processor means is responsive to the data transfer request for accessing the associated disc drive to obtain location data for the data and for generating a data message identifying the location, and the bus interface means is responsive to the data message for outputting the data message to the bus means.

8. A system as claimed in claims 6 and 7, in which each of the control means is responsive to the data message being outputted to the bus means for reading the data message from the bus means, and the control means in response to a data message identifying a file stored on its associated disc pack retrieves the data and outputs the data to the bus means.

9. A system as claimed in claim 8, in which the requesting one of the processor means is responsive to the data appearing on the bus means for enabling the bus interface means to remove the data from the bus means, and the control means further comprises a processor interface means responsive to the data for outputting the data to the requesting processor.

10. A system as claimed in claim 4, in which each of the control means comprises processor means which is responsive to a data request applied to the control means by one of said one or more processors for identifying the disc drive containing the data, and a disc controller which includes bus interface means for receiving a request via the bus means and which is responsive to the identification of the disc controller associated with the disc drive containing the data for transferring the data to the requesting processor.

11. A disc intercommunication system substantially as hereinbefore described with reference to Figs. 2 and 3 of the accompanying drawings.

Printed for Her Majesty's Stationery Office
by Burgess & Son (Abingdon) Ltd.—1982
Published at The Patent Office, 25 Southampton Buildings,
London, WC2A 1AY, from which copies may be obtained.